

AD-A134 275


MODELING EFFORTS FOR THE AUTODIN II FLY-OFF STUDY(U)
DEFENSE COMMUNICATIONS ENGINEERING CENTER RESTON VA
R PRISHIVALKO AUG 83 DCEC-TN-8-83


1/1

UNCLASSIFIED

F/G 9/2

NL





MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

AD-A134275

TN 8-83




DEFENSE COMMUNICATIONS ENGINEERING CENTER

TECHNICAL NOTE NO.8-83

MODELING EFFORTS FOR THE AUTODIN II
FLY-OFF STUDY

DTIC FILE COPY

AUGUST 1983



Approved for Public Release; Distribution Unlimited

83 11 01 034

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER DCEC TN 8-83	2. GOVT ACCESSION NO. AD-P-1577	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) MODELING EFFORTS FOR THE AUTODIN II "FLY-OFF" STUDY		5. TYPE OF REPORT & PERIOD COVERED TECHNICAL NOTE
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) R. PRISHIVALKO		8. CONTRACT OR GRANT NUMBER(s)
9. PERFORMING ORGANIZATION NAME AND ADDRESS Defense Communications Engineering Center Systems Analysis Division, R820 1860 Wiehle Avenue, Reston, VA 22090		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS N/A
11. CONTROLLING OFFICE NAME AND ADDRESS Same as Item 9		12. REPORT DATE August 1983
		13. NUMBER OF PAGES 38
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES Review relevance 5 years from submission date.		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) AUTODIN II, PSN, IMP, host, packet, transaction, access area, backbone net, TAC, TCP/IP, topological network design, AUTODIN II "fly-off"		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) The Systems Analysis Division, R820, provided support to the design teams and evaluation team of the AUTODIN II "fly-off." This support included collection and refinement of requirements, modification and application of computer models, verification of the proposed designs, and evaluation of network survivability. This document describes the procedures, computer models, and assumptions used, as well as the problems encountered.		

DD FORM 1 JAN 73 1473

EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A/1	



TECHNICAL NOTE NO. 8-83

MODELING EFFORTS FOR THE AUTODIN II
"FLY-OFF" STUDY

AUGUST 1983

Prepared by:
• R. Prishivalko

Approved for Publication:

GERALD W. MUSSELMAN
Colonel, USAF
Deputy Director
Engineering Applications and Analysis

FOREWORD

The Defense Communications Engineering Center (DCEC) Technical Notes (TN's) are published to inform interested members of the defense community regarding technical activities of the Center, completed and in progress. They are intended to stimulate thinking and encourage information exchange; but they do not represent an approved position or policy of DCEC, and should not be used as authoritative guidance for related planning and/or further action.

Comments or technical inquiries concerning this document are welcome, and should be directed to:

Director
Defense Communications Engineering Center
1860 Wiehle Avenue
Reston, Virginia 22090

EXECUTIVE SUMMARY

During the AUTODIN II "fly-off", the Systems Analysis Division, R820, provided support to the design teams and evaluation team. This support included collection and refinement of requirements, modification and application of computer models, verification of the proposed designs, and evaluation of network survivability. This document describes the procedures, computer models, and assumptions used, as well as the problems encountered.

The state of the User Requirements Data Base presented a major initial obstacle to beginning the process of network design. The network design process was hampered by attempting to rapidly apply existing design tools suitable for the AUTODIN II approach to the different Replica approach. This document describes the status of the design process during the "fly-off" period (i.e., through January 1982). Significant changes to this process will have occurred before publication of this technical note.

TABLE OF CONTENTS

	<u>Page</u>
EXECUTIVE SUMMARY	iii
I. INTRODUCTION	1
II. REQUIREMENTS	2
1. User Requirements Data Base (URDB)	2
2. Problems with the URDB	2
3. User Requirements Methodology	7
4. The Traffic Model	8
a. Overview	8
b. Non-Mode 6 Terminal-to-Host Traffic	9
c. Remote Concentrator and Mode 6 Terminal-to-Host Traffic	10
d. Host-to-Host Traffic	10
e. Traffic Model Implementation	11
5. The Overhead Model	11
III. DESIGN PROCESS	15
1. Introduction	15
2. Access Area Design	17
a. Switch Site Selection - the GUPTA Program	17
b. Access Line Costing - CULINES Programs	18
c. Other Available Access Design Programs	18
3. Backbone Area Design - GRINDER	19
a. Background	19
b. DCEC GRINDER - General Capabilities	20
4. Costing of Transmission	21
a. Current Practice	21
b. Effects of Tariffs	21

TABLE OF CONTENTS (Continued)

	<u>Page</u>
c. Overseas Costing	22
5. Retrospective Overview	22
IV. VERIFICATION	24
1. Introduction	24
2. Access Area Costing	24
3. Backbone Network Verification	25
4. Lack of Performance Criteria	25
V. SURVIVABILITY EVALUATION	27
1. Survivability Model	27
2. Application of Model	27
VI. CONCLUSIONS	30
ACRONYMS	31

LIST OF FIGURES

<u>Figure</u>	<u>Title</u>	<u>Page</u>
1	Network Cost Curves	16
2	Sample Drawdown Curve	28

LIST OF TABLES

<u>Table</u>	<u>Title</u>	<u>Page</u>
I.	Transmit and Receive (TAR) Dataset	3
II.	Trunks Dataset Format	5
III.	Replica Network Overhead	13
IV.	AUTODIN II Overhead	14

I. INTRODUCTION

In September 1981, the Director, DCA, General Hilsman, directed that a study be done to reassess the direction of the AUTODIN II planning, which would culminate in a Defense Data Network (DDN). Two design teams were established, each of which was to propose a DDN alternative. One team would use the existing AUTODIN II plans as a starting point, the other an approach based on ARPANET/WIN (WWMCCS Interconnect Network) concepts. Ground rules were developed for the two designs, and a third team was selected to evaluate the two alternatives once they were specified. The competitive nature of the study resulted in its being referred to as the DIN II "fly-off". The design teams were the DIN II team, which extended previous AUTODIN II plans in consonance with Western Union, and the Replica Team, which proceeded from technology and concepts previously applied in the ARPANET and WIN.

The Defense Communications Engineering Center (DCEC) provided support to both design teams and the evaluation team. In particular, the Systems Analysis Division, R820, collected and refined basic network requirements; modified, explained, and applied computer models used in the alternative network designs; verified the proposed designs by use of computer models; and evaluated the network survivability offered by the two DDN alternatives.

This document describes the efforts undertaken in these areas, and provides an overview of some of the problems encountered during the study and how they were resolved, as well as a description of the basic computer models used and their capabilities. It should provide useful information for anyone interested in the overall data network design process, in particular as applied by DCEC to the DIN II "fly-off".

II. REQUIREMENTS

1. USER REQUIREMENTS DATA BASE (URDB)

AUTODIN II was specified as a common user digital communications network that will provide subscribers (users) with data communication services for interactive timesharing and transaction-oriented systems requiring rapid response between terminals and computers, and for remote job entry and computer-to-computer data transfers requiring high speed transmission. The AUTODIN II User Requirements Data Base (URDB) was established to provide a central source of information about the user requirements that the network must satisfy. This user information consists of user locations, connectivity, interconnection technical requirements, traffic volume, and the line speed required to access the common user network. The URDB data were accumulated by asking the services and DoD to provide this user information about all operational and planned ADP systems and data networks that require long haul and area data communications support. The systems information was submitted by the agencies in the form of 80 column data cards, of which there are seven kinds. To facilitate the processing of this data for topological design purposes, certain information was extracted and formatted to produce a Transmit and Receive Traffic (TAR) dataset and a Trunks dataset. The TAR dataset contains basic information about the components of each ADP system; the Trunks dataset indicates the logical flow of traffic between the components of a system. (See Tables I and II for the TAR and Trunks dataset formats.)

2. PROBLEMS WITH THE URDB

Unfortunately, as with many data bases, an accumulation of information placed in a predefined format results in at best a collection of "raw" data. Frequently in a design exercise the analyst/designer discovers his initial, often significant, obstacle is to produce a usable, complete, and consistent set of requirements data, even though a data base already exists. The analyst is often the first individual to try to use stored data in a comprehensive fashion. Past AUTODIN II studies within DCEC had already revealed numerous errors and inconsistencies within the URDB. The analyst or data base user is, however, rarely in a position to directly correct the official data base. Furthermore, when inconsistencies are found, the incorrect data entry is not always evident. Nevertheless, for a meaningful design, consistency must be obtained.

More specifically, the errors present in the URDB were of several types. An ADP system as described in the URDB might not correspond with the real system. That is, the number and location of hosts and terminals disagreed with reliable external sources of information on that existing system. For example, one ADP system was described as having numerous terminals but no host computer, and the description of WIN differed radically from the true computer locations. The degree of preciseness of system descriptions varied widely, depending on how thorough the original providers of the data were.

The TAR and Trunks datasets produced directly from the URDB are merely extracts of some of the URDB entries, and therefore are no more error-free than their source. The TAR/Trunk process does, however, attempt to add

TABLE I. TRANSMIT AND RECEIVE (TAR) DATASET

<u>Field/Position</u>	<u>URDB Source/Comments</u>
Location (2-17)	Card 10, field 3 if military Card 10, field 4 if in a city
State Code (19-20)	Card 10, field 5
System Code (31-33)	Card 10, field 1
Latitude (35-41)	Latitude in degrees, minutes, seconds of location
Longitude (43-50)	Longitude in degrees, minutes, seconds of location
Location Number (52-55)	Unique Location Code
System Type (57)	Star Topology = 5, Non-Star Topology = 4
Transmit Traffic (58-66)	Busy hour transmit traffic in kb/hr
Receive Traffic (67-75)	Busy hour receive traffic in kb/hr
No. of Lines (76)	Blank implies 1
Linespeed (77)	Card 20, field 9 for terminals, Card 31, field 10 for hosts/remotes
Dual Home (79)	*if dual homing required
Star Host (81)	*if the host of a star topology system
Subscriber ID (82-85)	Card 10, field 2
Device Type (87-96)	Terminals - Card 20, field 3 Hosts/Remotes - Card 31, field 3

TABLE I. Continued

<u>Field/Position</u>	<u>URDB Source/Comments</u>
Mode (98-100)	Hosts/Remotes - set to '6', Terminals - Card 20, field 11
Area Code (102-104)	Card 10, field 7
Local Exchange (106-108)	Card 10, field 8
Cutover Date (110-113)	In Service Date
DU/FP (120)	Dial Up or Full Period

TABLE II. TRUNKS DATASET FORMAT

<u>Field (position)</u>	<u>Source/Comments</u>
LOCATION 1 (2-17)	Card 10, field 3 if on military installation Card 10, field 4 if in a city
STATE CODE 1 (19-20)	Card 10, field 5 for location 1
HTR 1 (22)	The first character of device identification at location 1 (H, R, or T).
LATITUDE 1 (24-30)	Geographical latitude of location 1
LONGITUDE 1 (32-39)	Geographical longitude of location 1
LOCATION 2 (41-56)	Card 10, field 3 if on military installation Card 10, field 4 if in a city
STATE CODE 2 (58-59)	Card 10, field 5 for location 2.
HTR 2 (61)	The first character of device identification at location 2 (H, R or T)
LATITUDE 2 (63-69)	Geographical latitude of location 2.
LONGITUDE 2 (71-78)	Geographical longitude of location 2.
SYSTEM CODE (80-82)	Card 10, field 1
SYSTEM TYPE (84)	Networks with a Star Topology = 5 Networks with a non-Star Topology = 4
TRAFFIC FROM 1 (85-93)	For the given link (Location 1 to Location 2) the Peak Busy hour traffic (card 60, field 9) on the transmit link direction (field 2) summed over each application type. Expressed in Kilobits/hr.
TRAFFIC TO 1 (94-102)	Same as TRAFFIC FROM 1 except substitute receive for transmit.

TABLE II. Continued

<u>Field (position)</u>	<u>Source/Comments</u>
LINESPEED (104)	Host to terminals, use card 20, field 9. For host/remotes, use card 31, field 10. Lowest value of the two devices.
FP/DU (107)	If card 50, field 3 is 'FP', set to 'F'; otherwise set to 'D'.
LOC 1 ID (117-120)	The device identification for location 1.
LOC 2 ID (121-124)	The device identification for location 2.

geographical coordinates to the URDB location names, and its failure to do so in many cases is caused by another URDB problem - invalid or misspelled location names. For any network design, the analyst needs to know where the users are located. The URDB instructions direct that the location be entered by its standard DCA contraction code and that it be a geographical location. Nevertheless, many spellings may occur for, say, Wright Patterson, and entries such as "HQ 5th Battalion" are not unusual. Determining just the number of user locations and where they are requires tedious effort.

Inconsistency of a different kind may be found among the entries for another user requirement. One field is supposed to express the line speed required for the access line from the user to the common user net. Another field expresses the peak hour, or busy hour, traffic in kb/hr transmitted and received by the user. Frequently, the requested line speed cannot accommodate the projected traffic. Clearly, at least one entry is wrong in such a case, but which one? The designer needs access line information for access area design, and traffic data for backbone network design. Other line speed and traffic inconsistencies exist on the system level. The sum of all transmitted traffic over all system components should equal the sum of that received; the URDB data may not reflect this. The line speed required for the only host computer in a system with numerous terminals should be greater than any of the terminal line speeds. Again, the URDB may have suspicious entries for some systems.

The URDB and TAR datasets have a mode field which should indicate the type of terminal being described. DCEC engineers believe that the surprisingly high number of mode 6 (IBM Synchronous Data Link Control) terminals in the URDB is probably the result of data supplier misunderstandings rather than technical realities. Nevertheless, the designer/analyst cannot arbitrarily decide that any specific mode entry is incorrect.

Other more minor errors also exist. The net result is that, even for the ADP systems represented therein, the URDB proved to be far from the usable collection of clean user requirements that any worthwhile design effort needs. The DIN II "fly-off", with its tight schedules and high-level visibility, drew attention to the true status of the URDB and hopefully also drew renewed emphasis to the need for verification and correction of its information.

3. USER REQUIREMENTS METHODOLOGY

Past AUTODIN II studies faced similar problems with the URDB. TAR and Trunks datasets had been created, and then iteratively processed both manually and through several computer programs to reduce the number of errors. Personnel within the comptroller's office had even validated the majority of the ADP systems of interest to them during past studies by contacting the specific agencies responsible for the original data input. Many corrections were made to the TAR dataset extract.

This TAR dataset thus produced was clearly the most accurate set of data available for the ADP systems that it attempted to represent. It therefore became the primary source of system descriptions for the DIN II "fly-off".

Unfortunately, it did not contain all the systems nor all the specific data required.

This basic TAR dataset represented the CONUS portion of its ADP systems only. Past studies had focused on CONUS only, whereas the DIN II "fly-off" was to consider overseas requirements as well. Furthermore, the comptroller's office was concerned with data needed for system costing only. Therefore, their validated TAR dataset provided the most accurate system component identification, locations, and access linespeed, but no traffic volume data. Finally, new systems had been identified for inclusion into the DIN II requirements that previously had not been considered.

The following methodology was developed and applied to produce, from this incomplete base, the data required for the alternative networks design:

- o For those systems within the base TAR dataset, information was extracted from the URDB for their overseas components. These newly added records required iterative processing to remove even the most obvious errors.
- o New data were collected and entered for several ADP systems; i.e., military ARPANET, MINET, COINS, IDHSC, WIN, I-S/A AMPE, and SACDIN.
- o A traffic model was agreed upon and used to generate transmit and receive busy hour traffic volumes, and the distribution thereof, for those systems having no better source of traffic information available. This proved to be all systems other than SACDIN, and the I-S/A AMPE's.

The net result, although obviously not a precise picture of the user requirements, was at least a representative, reasonably consistent, and usable set of data that allowed the design process to be undertaken. Within this new TAR dataset, 91 ADP systems comprising 488 host computers, 87 remote concentrators, and 1359 terminals were described. Any assumptions used in generating this data were mutually agreed upon by the two competing design teams.

4. THE TRAFFIC MODEL

a. Overview. As noted earlier, traffic volume and distribution had to be generated for most ADP systems. For each record, there already existed a required access line speed, which serves as a key indicator of possible traffic volume. For the purposes of developing the traffic model, a number of assumptions were made which are described below.

For this model's purposes, traffic is assumed to flow between the components of each individual system, and each terminal communicates with just one host. The terminal characteristics also influence the traffic determinations. A keyboard terminal is viewed as receiving several times more data than it sends; a mode 6 terminal is deemed able to transmit data at higher rates, e.g., from a tape. A mode 2A terminal is considered to serve interactive needs; a mode 1B is viewed as a query/response terminal.

The distribution of the traffic flow within an ADP system depends upon the number of its hosts and terminals. In a star system, one host supports all the terminals. Therefore, the sum of the terminals' transmitted traffic equals the host's received traffic. Likewise, the terminals' received traffic should equal the host's transmitted traffic. In an all-host system, it is assumed that each host communicates with all the others, sending as much traffic as it receives, within the constraints of the access line utilization. In a system with several hosts and several terminals, each terminal is "homed" to just one host. If, after all communications between terminals and hosts are considered, host access lines are still not utilized to the assumed busy hour limit, then traffic flow between hosts is assumed in order to realize the host access line limit. The distribution of transaction types on host-to-host connections is the same as that described in the AUTODIN II specifications.

Seven message or transaction types are considered. The data portions of each type are assumed as follows:

<u>Message Type</u>	<u>Length (Bits)</u>
Inquiry	240
Response to Inquiry	480
Query	560
Response to Query	23000
Narrative	12000
Bulk	500000
AUTODIN I	27000

b. Non-mode 6 Terminal-to-Host Traffic. Terminals which are not mode 6 are divided into two classes: mode 2A and mode 1B. The traffic model assumes that mode 2A terminals send inquiries to and receive responses from the single host with which the terminal communicates. The rate at which these inquiry/response transactions occur during the busy hour is determined by the terminal speed, according to the following table:

<u>Terminal Speed (b/s)</u>	<u>Transactions/Min</u>
< 1200	5.8
≥ 1200	7.2

Mode 1B terminals are assumed to send queries to their host and receive responses to queries. The rate at which these exchanges occur during the busy hour is again dependent on the terminal speed, as follows:

<u>Terminal Speed (b/s)</u>	<u>Transactions/Min</u>
< 300	.2
≥ 300, < 1200	.4
≥ 1200	.7

For example, a mode 1B terminal shown as requiring a 2400 baud access line to the network is assumed to send .7 560-bit queries per minute on the average to its host during the busy hour, and receive .7 23,000-bit responses per minute. This implies a busy hour transmitted traffic volume of 23.5 kb/hr, and a busy hour received traffic of 966 kb/hr.

c. Remote Concentrator and Mode 6 Terminal-to-Host Traffic. Each remote concentrator and mode 6 terminal communicates with a single host. The concentrator and terminal generate and receive all seven types of messages, in a symmetric exchange with the host. The model assumes the busy hour data flow uses one-third of the device's access line capacity. If the mode of a remote concentrator is indicated as 1B, half-duplex operation is assumed, and thus, the busy hour data flow is assumed to be only one-sixth of the access line speed.

The distribution of the types of messages sent and received is given by the following table, where M is the total number of messages or transactions per second:

<u>Message Type</u>	<u>Message Rate</u>
Inquiry	.135M
Response to Inquiry	.135M
Query	.07M
Response to Query	.07M
Narrative	.26M
Bulk	.11M
AUTODIN I	.22M

Since the total data flow in bits per second is known, being directly derived from the device type, mode, and line speed, and the length in bits of each message type is known as well, M can be computed. With M known, the number of transactions of each type per busy hour can be found.

d. Host-to-Host Traffic. Host-to-host traffic is symmetric. If host A sends M messages to host B, then host B sends M messages of the same type to host A. Of course, two hosts communicate only if they are designated as part of the same ADP system. The process for computing the volume of host-to-host traffic is as follows.

Assume there are several hosts in a particular ADP system. Each host is assigned a maximum data flow in each direction, based on its required access line speed. If the line speed is ≤ 9600 b/s, this maximum is one-half the line speed; otherwise it is one-third the line speed. For each host, the data flow on its access line to and from each of its terminals or remote concentrators is computed, as already described. The remaining usable capacity on each host's access line is then computed. The host with the least (but greater than zero) spare capacity is then assumed to use the remainder of its allowable data flow to exchange messages equally with the other hosts in the system whose lines are not filled to their limits. This data flow is considered in further reducing all participants' spare capacity. The new minimum positive spare capacity host is determined, and another amount of data

flow is distributed and assigned between him and the remaining hosts having spare capacity on their lines. This increasing of traffic volume continues until either all host access lines are filled to their data flow limit, or only one host still has spare capacity. The message types exchanged host-to-host are assumed to be in the same proportions as in the mode 6 terminal-to-host model.

e. Traffic Model Implementation. The model was implemented to produce the transmit and receive busy hour traffic (in kilobits) for the requirements TAR dataset, and also to produce a Trunks dataset that reflected the logical connectivity defined by the model.

The first step was to generate the transmit and receive traffic for each terminal. This was done, as described earlier, by considering only the mode and line speed associated with each terminal. Next, to develop host traffic volumes, the homing of the terminals had to be known. An arbitrary rule was used to assign homings in the non-star systems. Each terminal was homed to the geographically closest host within the same ADP system, unless that host already had more than its proportionate share of terminals assigned. In that case, the second closest host was considered. This homing resulted in pseudo-star systems being defined. For either a real or pseudo-star system, the terminals' transmit traffic and receive traffic were summed to generate the host's receive and transmit traffic respectively. The star system traffic generation was thus completed.

For an all-host system, the host-to-host model was applied, as previously described.

For a hybrid system of hosts and terminals, the above process already provided all traffic for the terminal TAR records, and, through the pseudo-star system method, part of the traffic for the host TAR records. The host-to-host model was then applied, as for the all-host systems, to generate the remainder of the traffic for the hosts.

At each step of traffic assignment, trunk dataset records were also created to represent just how much traffic was assumed to be flowing between any pair of system components. The primary trunk dataset showed this information in kilobits per hour.

Message or transaction trunk datasets were also established. These indicated the assumed pairwise flow of transactions of each type per busy hour. They were created to allow for proper assignment of overhead bits to the data flows, so that, eventually, alternative-dependent traffic matrices with overhead could be used in backbone network design.

5. THE OVERHEAD MODEL

The alternative designs required a common set of user requirements in order to be meaningful and comparable. For most ADP systems, the traffic model described generated the requirements in terms of data bits and types of transactions that needed to be delivered by the designed common user network. However, some amount of traffic-dependent overhead, depending upon the

particular network and its implementation methods, must also be carried by the designed system. This overhead, both in terms of additional bits and additional packets, was provided by each design team.

Generally, the overhead specifications included consideration of segment overhead and packet overhead, both Transmission Control Protocol (TCP) overhead in the forward flow direction, and TCP acknowledgements (ACKS) in the reverse data flow direction. The practice of piggybacking TCP ACKS when possible was taken into account. For each type of transaction, the expected numbers of bits and packets added to the network in both forward and reverse flow directions were calculated. The maximum packet size of each alternative was another factor in the calculations. The results are given in Tables III and IV.

Each previously generated transaction trunk dataset was used with the appropriate overhead factors to generate, for each alternative design, the peculiar bit flows and packet flows that are required.

TABLE III. REPLICA NETWORK OVERHEAD

TRANSACTION TYPE	DATA BITS	FORWARD OVERHEAD			EXPECTED # PACKETS	REVERSE OVERHEAD		
		OVERHEAD BITS	RATIO: TOTAL BITS/DATA BITS	TOTAL BITS		OVERHEAD BITS	RATIO OH/DATA	EXPECTED # PACKETS
Inquiry	240	699.3	3.91		1.4	474.6	1.98	1.4
Response to Inquiry	480	699.3	2.46		1.4	474.6	.98	1.4
Query	560	699.3	2.25		1.4	474.6	.85	1.4
Response to Query	23000	6997.2	.32		26.6	1898.4	.08	5.6
Narrative	12000	3598.6	1.30		13.8	949.2	.08	2.8
Bulk	500000	137553.8	1.28		549.4	31323.6	.06	92.4
AUTODIN I	27000	7797.2	1.29		30.6	1894.4	.07	5.6
SACDIN	2000	3099.3	1.55		3.4	838.5	.42	1.4

TABLE IV. AUTODIN II OVERHEAD

TRANSACTION TYPE	FORWARD OVERHEAD				REVERSE OVERHEAD		
	DATA BITS	OVERHEAD BITS	RATIO: TOTAL BITS/DATA BITS	EXPECTED # PACKETS	OVERHEAD BITS	RATIO OH/DATA	EXPECTED # PACKETS
Inquiry	240	792	4.30	1.01	309	1.29	.4
Response to Inquiry	480	792	2.65	1.01	309	.64	.4
Query	560	792	2.41	1.01	309	.55	.4
Response to Query	23000	3983	1.17	6.01	1543	.07	.4
Narrative	12000	2360	1.20	3.01	929	.08	.4
Bulk	500000	86248	1.17	110.01	33938	.07	.4
AUTODIN I	27000	4712	1.18	6.01	1854	.07	.4
SACDIN	2000	792	1.40	1.01	309	.16	.4

III. DESIGN PROCESS

1. INTRODUCTION

The term "design" as used throughout this document should perhaps more properly be qualified as topological design. The goal of this design process is to define, at least topologically, the best network that satisfies the specified user requirements. Most often for DCEC studies, this further translates into defining the least expensive network that satisfies all specified constraints. The specified constraints, in addition to the user-to-user traffic demands, normally include security, survivability, and performance constraints on the network.

The topological design of a common user network is often split into an access area and a backbone area design. The access area portion may include selection of number and geographical placement of switches, selection of number and placement of concentrators and/or multiplexers, and access line layout from the user to the switch via any intermediate level devices.

The backbone area design includes determination of the backbone trunking layout and trunk sizing. The trunking decisions, of course, require the switch-to-switch traffic requirements which are developed from the user-to-user requirements and the homings of the access area design.

In a simplified view, the main variable cost factors are access transmission costs, switch costs, and backbone transmission costs. For a given number, N , of geographical switch locations, the access homings and transmission costs can be calculated, and then a backbone cost can be computed. In other words, for any N , the component costs and hence total costs for the lowest cost network satisfying the design requirements can be found. Generally, as long as cost increases with distance, as N increases, access transmission costs will diminish and backbone transmission costs will increase. The tradeoffs are represented by the graph in Figure 1. The design process may be alternatively viewed, therefore, as development of the tradeoff curves in the graph for the applicable requirements and tariffs, or at least development of enough of the curves so that the minimum point may be found on the total cost curve.

The comprehensive design process requires iterative executions of numerous subnetwork design and costing programs. The time constraints of the DIN II "fly-off" study prevented a thorough search for the overall best topology; fine tuning of the proposed network solutions no doubt could have been done if more time had been available. The DCEC role in the design process varied between the two design teams. For the DIN II team, DCEC served as a computer model provider, modifier, and instructor, generally supplying the tools for that team to undertake the design process. For the Replica team, DCEC applied several of these same models, particularly in pursuing the best access area design. Later, many of these same models were used in verifying the final proposed network designs of both teams.




Figure 1. Network Cost Curves

This section briefly explains the various computer models used during the DIN II "fly-off", and also how they were applied within DCEC for the Replica team.

2. ACCESS AREA DESIGN

a. Switch Site Selection - The GUPTA Program. The GUPTA computer program is a combinatorial model that aids in selecting the lowest cost access area design. The following problem is addressed by it. There is a set of user locations, from each of which some number of access lines must be homed to a higher level node (e.g., Terminal Access Controller (TAC), switch) of the network. Given a set of candidate locations for switch placement, and a specified number of locations, N , the GUPTA program attempts to select the best N locations to be used. Best is defined here as that set of N locations that allows the least expensive homing of all users. The name GUPTA refers to the person given credit for the heuristic algorithm used within the program to approach the combinatorially enormous problem of selecting the best N items from some larger set.

As input, the program requires the user locations, number of required lines and their speeds by location, the list of candidate sites, the appropriate costing routines, and specification of N . The outputs will include the selected N sites, the homings upon which this choice is based, and the associated costs.

The program, of course, does not address all access design questions. A variety of values for N must be used, and hence the program must be run iteratively in any quest for the best design. Furthermore, the program can address only one stage of any multilevel access design at a time. For example, if terminals must be homed to TAC locations, and then the chosen TAC's must be homed to switches, the GUPTA program must be set up and executed twice in order to represent this one choice of a number of TAC's and a number of switches. If the designer wishes to investigate 4 different numbers of TAC locations and 5 possible numbers of switch locations, 20 runs of GUPTA would be required to cost out all the possibilities. The program also does not consider capacity constraints on a switch or TAC. If a location already has n lines homed to it, and n lines is all that one switch can terminate, the program will proceed to consider homing yet another line there, without noting that the cost of a new switch would be incurred as well.

Despite the suboptimalities caused by the problems listed, the GUPTA program is still very useful, even if sometimes awkward to apply. For the DIN II proposed design, where terminals could be homed directly to the TAC collocated in the PSN (packet switching node), the multilevel problem did not present itself. For the Replica design, the situation was more complex.

The Replica network requires that all terminals be homed to TAC's, which should be located at the most cost-effective sites. Furthermore, for security reasons, a classified system terminal can be homed only to a TAC of the same classification. The result is that a set of unclassified TAC's must be located for the unclassified terminals, and likewise separate sets of TAC's for the Confidential, Secret, Top Secret, and Special Intelligence classifications. Thereafter, a best set of switch sites must be selected to

which all TAC's and hosts are homed. For any N switch sites, the choice of best sites and the total access cost depend on the particular lower level decisions of TAC locations. The number of possible combinations is unmanageable.

Fortunately, in practice, rational decisions can greatly reduce the number of needed GUPTA runs by narrowing the possibly good choices. For example, if a classified TAC is expensive, and the number of terminals of that classification is small, clearly a small number of TAC's is best, since savings in transmission costs are unlikely to compensate for the TAC costs. Furthermore, bounds on the potential cost savings can usually be found by running extreme cases (e.g., one TAC versus a TAC for each terminal location), and the degree of investigation into the best answers may be strongly guided by the likely magnitude of any savings. The expense of hours or days of man and computer time is probably not justified in order to "save" \$10,000 in access cost in a multimillion dollar network.

Therefore, the approach of the Replica team in using GUPTA to address access design was to consider two or three choices of number of locations for each TAC classification level. For the small subscriber sets (i.e., TS, SI terminals), the better number of locations was obvious quickly. Overall, after perhaps 25 GUPTA runs, a reasonable access area design and cost for each of 60, 80, and 100 switch sites could be obtained.

b. Access Line Costing - CULINES Programs. During past AUTODIN II studies, a large number of costing programs were written. One basic program is called CULINES (Common User LINES). Given a predetermined set of switch locations as input, the program will determine the transmission cost of connecting each subscriber to his best switch with the appropriate required access line. The program does output homing information in that it determines to which switch a subscriber should be connected to minimize his access cost. Its primary utility is that it is organized to interface with a TAR dataset, and to maintain ADP system identifiers in outputting of costs. It will also indicate the cost savings possible through multiplexing access lines originating from the same geographical location. It will cost dial-up as well as full-period lines.

CULINES, however, has limited utility as a design aid. As already stated, the set of switches to be used must be prespecified. If similar costing is used, the homings produced by CULINES are identical to those produced by GUPTA. In the DIN II "fly-off", CULINES was used to provide a standard costing for both alternatives, and was also used because of its preestablished interfaces with other costing and traffic matrix generation programs.

c. Other Available Access Design Programs. Time constraints prevented the introduction of other models into the design process that would have been beneficial. One model that was not ready for use would have aided in more efficiently locating concentrators (e.g., TACS) and multiplexers. The models that are likely to be applied to similar, future data network design are discussed in the following paragraphs.

The awkwardness of the TAC location process for a Replica design was indicated above. Also, the only multiplexing considered in either GUPTA or CULINES was restricted to one geographical location at a time. A design model available within DCEC, known as MUXLOC (multiplexer locator) or CONLOC (concentrator locator), could better address these points, and therefore the model is briefly described.

The MUXLOC model uses what is known as an ADD algorithm. Assume the set of switches has been selected, and therefore the cost of direct homing the users is known. A set of candidate locations for concentrators (or multiplexers) is provided as input to the program. The program will find the best location to place a concentrator/multiplexer by computing the possible cost savings from rehomeing users through this device. The maximum capacity of any one concentrator/multiplexer is considered. The program will continue to add new concentrator/multiplexer locations one at a time until no more cost savings are possible. Savings possible by merging lines from several geographical locations into one device are also considered. The best total number of concentrators is found in one program execution, rather than by iterative GUPTA runs.

An algorithm approach can also be used to approximate the optimum number of switch locations to have in a network. This program (available in DCEC under program name ADD) begins with a zero backbone trunking cost and all users homed to one central switch. At each iteration within the program, the next best switch location is added to this developing network, as long as cost savings result. The considered tradeoff is between decreased access transmission costs and increased backbone transmission plus switch costs. Unfortunately, the incremental backbone trunking costs incurred by adding the N+1 switch to a network of N switches is not often known and may be difficult to approximate, especially for a totally new network. It should be emphasized that the MUXLOC and ADD programs, as well as GUPTA, are heuristic approaches. None can guarantee the optimal design. They all, however, provide insight into the access area design tradeoffs, and when used iteratively, or possibly together, can produce a cost-effective access topology.

3. BACKBONE AREA DESIGN - GRINDER

a. Background. The analytic model GRINDER (Graphical Interactive Network Designer) is the primary model used within DCEC for the topological layout and performance evaluation of a packet switching network backbone. GRINDER was developed by Network Analysis Corporation (NAC) in the mid-1970's. It is a comprehensive package written mostly in FORTRAN.

The original package was established to run on a Digital Equipment Corporation PDP-10, and was accessed by DCA users via the ARPANET. The graphical display features required utilization of an IMLAC terminal. The complete package addresses both access area and backbone area design.

To increase the ease of use of the model, DCEC analysts converted a number of modules in the original GRINDER package to code compatibility with an IBM 370 and Tektronix terminals. The resulting package, known as the DCEC GRINDER, contained only backbone area programs. Personnel within DCEC made

further modifications to this GRINDER, and incorporated several new analysis capabilities, to produce the DCEC GRINDER that was the model used for the DIN II fly-off. Several of the original GRINDER access area programs served as the basis for the MUXLOC and ADD programs described earlier.

The original (ARPANET) GRINDER continued to be used by personnel within DCA Headquarters for various tasks associated with the maintenance of the ARPANET. In June 1980, this group agreed to lease from the Network Analysis Corporation a slightly updated GRINDER package that was modified to handle more backbone nodes and to permit use of a Tektronix terminal for graphical display. NAC has restricted the dissemination of this package to DCA and DCEC only, and designated the program modifications as proprietary.

The net result is that there are, in essence, two GRINDER models. The DCEC version is a backbone-only model that is run on the DCEC ITEL-AS/5 computer. The ARPANET version is an access and backbone model, run remotely via the ARPANET on PDP computers, and subject to certain restrictions due to proprietary code. The remainder of this document will use the name GRINDER to refer to the DCEC version.

b. DCEC GRINDER - General Capabilities. GRINDER is an interactive backbone design and performance model. For the most part, it is an analytic model, rather than a simulation, that combines a number of algorithms which are based on viewing a packet switching network as a network of queues. GRINDER features an easily learned hierarchical command structure which aids the user in invoking the numerous capabilities.

Briefly, the model has the capability to:

- o Design basic network topology from an input of switch sites, switch-to-switch traffic, and an acceptable delay constraint;
- o Evaluate network performance under traffic routings that vary from strictly minimum-hop to high degrees of load splitting;
- o Output a variety of plots and statistics at each phase to guide the user's actions and decisions;
- o Iteratively invoke algorithms to modify a network to improve performance or lower cost;
- o Add, delete, or modify user-designated links;
- o Compute overall network reliability given link or node probabilities of failure.

Two important input parameters are the acceptable average end-to-end (originating switch-to-terminating switch) delay and the average packet length. In evaluating the performance of a specified, traffic loaded network, the model has the ability to compute the expected delays for packets of other than average size and/or priority.

GRINDER is still apparently a state-of-the-art tool, and its accuracy has been verified against actual ARPANET performance. Note that the computer resource requirements of GRINDER grow rapidly as the number of switch sites increases. This can cause some difficulties with its interactive use. A process such as initial design and refinement of a fairly large network (e.g., 70 switch sites) may require a large number of iterations of involved algorithms, with each iteration consuming over a minute of CPU time. Batch processing can be used for some of these steps, with the resultant network stored and available for later input for interactive analysis.

4. COSTING OF TRANSMISSION

a. Current Practice. The past AUTODIN II studies established a procedure for costing, as exhibited in the CULINES program, based upon a rate-step table. The rate-step table was simply a dataset of prespecified format which, for each of nine line speeds, indicated the fixed, per mile, subscriber modem and other charges used in calculating the cost of a transmission line. The per-mile charge changed at specified mileage break-points. Two rate structures were included: one to represent the ATT 260 series of tariffs, and the other the DDS tariff. The comptroller's office at DCA Headquarters created and updated the rate-step information as needed.

As previously stated, the purpose of the costing (as in CULINES) was to provide the final costing of both design alternatives in the DIN II "fly-off". For the sake of consistency, the same costing procedures were adopted for both designs. GUPTA and GRINDER were modified to read the rate-step table and to use the data, as provided by the comptroller's office, in the same way as CULINES.

Despite the many numbers present in a rate-step table, it really represents a very simplified view of the two tariffs. The procedure basically assumes that only three pieces of information are required to cost a transmission line. These three are whether the line is to be full period or dial-up, the line speed or capacity in kilobits per second, and the distance between the termination points of the line. Given these data, the cost is calculated for each of the two tariffs, with exorbitantly high figures being entered if the particular tariff is considered nonapplicable for the specified line speed. The lower of the two resulting cost figures is selected as the cost of the transmission line and any modems.

b. Effects of Tariffs. A major design goal is to minimize cost. The cost of the transmission lines, both access and backbone, is of course dependent on the assumed tariffs that are applied to the design. The simplified costing process used in the DIN II "fly-off", however, may not be appropriate for some design studies.

In reality, AT&T tariffs determine transmission charges for several line speeds, not merely by line speed and mileage but also by what rate centers the locations connected are in (A or B). The DDS tariff, deemed applicable to line speeds of 9.6 kb/s and greater, is currently available at only certain locations in the United States. In short, to accurately reflect currently available tariffs, the geographical location of the subscriber should be

considered. This added criterion will not only affect the total cost of a design, but also conceivably alter switch site selections, access line homings, backbone trunking, and other design specifics chosen by cost tradeoffs. On the other hand, certain details of current tariff charges may be ignored if the designed network is to be implemented years in the future when these details are likely to change. The general structure of the tariff (e.g., linearly proportional with distance, "breakpoint" tariff) is the most important tariff feature in long range planning.

As an example, access area homings of hosts to switches are usually determined by least cost. As long as the tariff charges increase with distance, each host is "best" homed to its closest switch (disregarding such things as switch termination constraints and survivability considerations). However, if the extra detail of rate center category (A or B) were considered, it is quite conceivable that some hosts would avoid their closest switch if it were in a B rate center. Backbone trunking is also affected by tariff structure. A breakpoint tariff with higher charges proportionately for shorter trunks would favor a design with more long haul trunking. Also, the amount of tandeming that is cost effective is influenced by the ratio of fixed charges to per-mile charges. The choice of connecting two switches with two tandem trunks versus one direct trunk may be determined primarily by the termination charges at the intermediate switch.

c. Overseas Costing. The design process described in the preceding sections was not used to design the worldwide network in one fell swoop. It was basically the process used for the CONUS portion of the network. The European and Pacific theaters were handled separately; differences in applicable costing were the primary reasons.

To keep the process as smooth as possible, the comptroller's office developed two additional datasets in a rate-step structure. One represented their approximation of costs to be used for Europe. The other was to be used for transoceanic connections, which would include most interisland hops in the Pacific. Numerous assumptions, of course, were required to generate these reasonably simple representations of costing in areas where international borders must be crossed and transmission is provided by several different common carriers as well as U.S. Government owned equipment. Specific information about the cost components used for these two datasets, and for the CONUS rate-step set as well, can be provided by the comptroller's office.

The cost and design models can use only one rate-step dataset at a time. Therefore, the users and switches had to be partitioned into CONUS, European, and Transoceanic subsets for each access and backbone area costing, for each design. However, once a worldwide network has been designed, the performance of the entire network may be verified all together, since the performance is independent of cost.

5. RETROSPECTIVE OVERVIEW

In general, the design process applied in the DIN II "fly-off" was significantly shaped by expediency. A conglomeration of programs existed at the start of the DIN II "fly-off" that had been written over the years and

modified hurriedly when necessary over the course of numerous past AUTODIN II studies. The programs varied in function from costing network pieces, creating reports, addressing design, and building a backbone traffic matrix, to data formatting and more. Although awkward to use, they were already established with the proper interfaces to the requirements datasets and the major design models. Furthermore, the reports produced were already familiar to the comptroller's office and other involved personnel. There was not time to create new programs to streamline the necessary steps. The DIN II "fly-off" did reveal, however, that the past programs were primarily packaged together to address the earlier AUTODIN II design questions. When different concepts were introduced, as in the Replica design, the existing set of programs proved to be incomplete, and in need of modification. Extensive analyst and programmer effort was required to successfully apply them to the DIN II "fly-off" study. Also note that the DIN II "fly-off" had a feature not present in many design efforts; it was a contest, more or less, between two teams. Whereas in any design process the analyst is frequently forced to make "reasonable assumptions" in order to allow work to continue, in the "fly-off" every assumption had to be jointly approved and reviewed for any possible unfair impacts on one team or the other.

One lesson that can be relearned from all this is that no well-done network design should ever be thought of as a "pushbutton" operation, no matter how many helpful computer models already exist.

IV. VERIFICATION

1. INTRODUCTION

Two alternative designs were proposed by the DIN II "fly-off" design teams. The stated design attributes had to be verified by a neutral party. DCEC was involved in verifying the performance and cost of each network's backbone, and also the cost of the access transmission lines used in the Replica approach. In order to permit this verification, each team, for its design, had to provide the node (i.e., TAC, IMP, PSN) sites selected, their homing strategy, and the backbone trunking details.

2. ACCESS AREA COSTING

As stated earlier, programs had already been established during previous studies and used by both DCEC analysts and comptroller's office personnel to cost an access line as described in a TAR dataset record. Multiplexing savings on a geographical location name basis were also considered. These programs were to serve as the common costing tools for both designs.

The Replica network had several characteristics that made costing of its access area transmission less than straightforward using these tools. The CULINES program, previously described, is both a costing and a homing program. Therefore, only that set of users and the TACS or switches to which they could be homed could be considered in one execution. Furthermore, only one rate-step dataset and only one level of a multilevel access area could be considered at a time. These restrictions resulted in a multitude of partitionings of datasets and a tedious costing process.

More specifically, because of the three different cost tables or rate-step datasets, any costing process had to be done separately for CONUS, Europe, and the Pacific. The costing process consisted of numerous steps. As previously noted, terminal to TAC homings had to be done separately for each of four security levels. Thereafter, the required access line speed to an IMP had to be determined for each TAC by considering the traffic from all the users homed to it. This new access line requirement had to be expressed as a TAR format record. With all TAC records determined, the TAC-to-IMP and host-to-IMP transmission lines could be costed. Although later changes in guidance from NSA altered the required homings at this level, at the time of the design proposal and verification, there was a perceived requirement for "dirty" and "clean" IMP's. Those costs considered related to C²I needs were allowed to be applied to a set of "clean" IMP's. Unclassified TAC's and other hosts had to be homed to "dirty" IMP's (i.e., IMP's other than those used for C²I hosts). Classified TAC's were permitted to be homed to any of the IMP's. In other words, further division was necessary in the costing process. In actuality, a very large amount of dataset manipulation and over a hundred program executions were used to perform all the steps required to cost the access transmission lines for the proposed Replica design.

3. BACKBONE NETWORK VERIFICATION

The backbone portion of each proposed network was cost and performance verified. Each team provided DCEC with a specification of its switch location and trunking.

The costing process was simple. The GRINDER model had been modified to use, in all its costing, procedures identical to those in CULINES. The CONUS and European trunk costs were computed in separate runs using the appropriate rate-step dataset. Transoceanic trunks, very few in number, were costed manually.

Performance measures were computed for the worldwide network as a whole. A primary input to this process was a switch-to-switch traffic matrix, constructed from the user homings determined during the access area design. Each traffic matrix included both the data bits and the appropriate traffic-dependent overhead bits for the respective system.

For each proposed design, the following performance verification process was used. GRINDER routed the switch-to-switch traffic over the prespecified network, using two iterations of its routing algorithm (i.e., using some load splitting). Delays to be expected on a per-packet basis over a network thus loaded were then computed.

Delay calculations were made for the average size packet of the particular network. The underlying assumption used in these delay figures was that all traffic was of the same priority. The calculated delays indicated the expected end-to-end (originating switch to terminating switch) delay for the specified packet for each pair of switches in the backbone. The worst cases, or maximum expected delays, were noted. The cumulative delay distribution could also be calculated for any pair of switches. The expected average delay was often several times less than the 99th percentile delay.

Similar delay calculations were done for packets of 600 bits in size, which were assumed to be 1 percent of the total number of packets, and of top priority. Performance constraints had been set for these important packets across the backbone.

The routing of the traffic also resulted in statistics showing the number of kilobits expected to tandem through each switch. The average packet size was used to convert these numbers to an expected number of tandem packets. The number of tandem packets was added to the number of originating and terminating packets at each switch, available from the specified user homings and transaction type assumptions, to compute the total number of packets requiring processing at each switch site. This was used to verify that the number of IMP's/PSN's proposed for each site could handle the expected packet load.

4. LACK OF PERFORMANCE CRITERIA

The preceding section discusses the performance measures that were calculated for each design's backbone network. In order to evaluate the

performance of any design, and certainly to compare two different ones, standards of some sort must be set, such as an acceptable level of performance against which each design can be judged. An adequate set of performance criteria did not exist in the DIN II "fly-off".

The only performance constraint clearly defined for both designs was that a Category I packet (i.e., highest priority, see System Performance Specification for AUTODIN II, Phase I), with an assumed size of 600 bits, be able to traverse the backbone between any switch pair in 200 milliseconds. No minimum delay was agreed to for the remaining 99% of the traffic, with the result that the numerous performance outputs of the modeling programs could not easily be used for evaluation purposes. Networks could be designed that would satisfy the high priority 1% of the traffic of small packet size, and yet give abominable service to all other packets. No meaningful conclusions can be drawn from the delay calculations for the average size packets of each network. The average packet size in one network was 996 bits, in the other 4499 bits. In other words, the information content of the "average" packet differs greatly between the two networks, so a direct comparison of "average" packet delays provides no immediate insight. Moreover, the important criterion in any network is the delay time to complete an entire transaction from end user to end user. This may mean sending one packet, or several packets, depending on the transaction size and system packetizing procedures, across the backbone, as well as transmitting the transaction intact or piecemeal through the access area. The proportion of the entire delay experienced over the backbone will not be constant between the two different designs.

To summarize, to permit meaningful performance comparisons, acceptable user-to-user delays should be determined for each type of transaction. Analysis should be done to estimate the delays experienced throughout the entire network, including the access area and during the segmentation/packetizing process. A model such as GRINDER is capable of providing the backbone delays on a packet basis. Design verification would then include verification that the acceptable transaction delays are met.

V. SURVIVABILITY EVALUATION

1. SURVIVABILITY MODEL

Comparison of the survivability of the two network backbones was another goal of the evaluation process. A computer model was used to quantify survivability in an overall network sense. The term robustness rather than survivability may be appropriate. The model addresses the question: How well can the network (backbone) continue to function as nodes are destroyed?

One assumption is that a measure of the quality of any backbone network is the number of terminal-to-host and host-to-host pairs that can communicate over it. "Communicate" in this sense simply means that a path over the operating links can be found between the switch serving one terminal or host and the switch serving the other host. Ideally, specific needlines could be identified as critical and only these considered, but the two design teams had considered different sets important. Therefore, all terminal-to-host and host-to-host communications were considered, with host-to-host pairs weighted more heavily. A more subtle underlying assumption is that the network routing algorithm can find any existing path.

The model also assumes a worst case analysis. At each iteration, one switch site still in the network is considered destroyed, and the resultant network is evaluated using the measure defined. The site "destroyed" is that which, of all the remaining sites, will most diminish the network measure. Furthermore, if a total of N sites is assumed destroyed, the model attempts to select the set of N sites that together cause the most damage. This is done by a heuristic algorithm similar to that in the program GUPTA. After site N is selected, the preceding N-1 choices are reconsidered, and changes made if any other site when considered with site N can lower the network measure more than one of the N-1 sites. This worst case analysis is graphically represented by a drawdown curve (see Figure 2). The same methodology can be extended to include transmission sites as well as switch sites. If a transmission facility supports a number of links, its outage can conceivably damage the network more than that of a switch site.

2. APPLICATION OF MODEL

The basic model was applied to the CONUS portion of each design team's backbone network. The primary output for each design was a drawdown curve, showing how rapidly the network measure could decline in the worst case analysis. The network measure is based upon the immediate impact of destroying a site, and does not consider contingency plans or reconstitution.

The number of terminal-to-host and host-to-host connections can be altered in two ways. One is by causing outage of backbone links. The other, more direct, way is by causing outage of the switch site to which the user is homed. (Dual homing, used in each design for only a small number of the total users, was not considered.) Destroying a switch site thus impacts the measure in both ways. Another interesting measure, therefore, is the percentage of user pairs homed to an operating switch that are still connected via the network. If a user can be easily disconnected from the majority of the




Figure 2. Sample Drawdown Curve

network without destroying him or his switch directly, then the robustness of the network is questionable.

The basic model considers each geographical switch site as distinct. A modification was made to address the issue of damage if the switch sites were targeted for nuclear attack. This was represented by assuming all switch sites within a 6 mile radius of the targeted site would be destroyed along with the target. The effect of this approach was to produce drawdown curves for each of the two designed backbones that could more meaningfully be compared.

VI. CONCLUSIONS

This document describes the modeling efforts undertaken within DCEC in support of the AUTODIN II "fly-off" study, in particular the computer models used and how they were applied during the various phases of the work. Specific results from the models are not presented herein.

Several problems encountered are also described. A brief recount of the major ones will serve as concluding remarks.

Attempts to apply the information within the User Requirements Data Base (URDB) have indicated a large number of errors in data accuracy and consistency. The data in the URDB need to be validated, and procedures should be established to ensure future entries are accurate and standardized before being permitted. In this way, the time-consuming process and numerous assumptions used in the DIN II "fly-off" to generate a consistent set of requirements may be avoided.

Any set of design tools is likely to require modification when applied to a new task. This statement is especially applicable if the programs were established with a specific system in mind, and the new task addresses a different system. If there is no time to modify the programs and streamline the process, the analyst is forced into an awkward procedure of making minimal essential changes only, and then writing and applying "quick and dirty" program interfaces in order to produce results.

The task of designing and comparing two alternative networks requires extra caution in making assumptions and approximations. The concern is that neither design be affected more than the other. This concern can slow the design process if parties of the different teams must agree at each such step; however, agreement is necessary.

Performance criteria were never specified to a degree sufficient to permit meaningful performance comparison of the two designs.

Hopefully, the model descriptions and problems presented can be instructive to anyone undertaking similar efforts in the future.

ACRONYMS

ADD	DCEC Program to determine optimum number of switches
ARPANET	Advanced Research Projects Agency Network
CONLOC	CONcentrator LOCator
CULINES	Common User LINES
DDN	Defense Data Network
DIN	AUTODIN (AUTOMATIC DIGITAL Network)
I-S/A AMPE	Inter-Service/Agency Automated Message Processing Equipment
IMP	Interface Message Processor
MUXLOC	Multiplexer LOCator
TAC	Terminal Access Controller
TAR	Transmit and Receive Traffic
TCP	Transmission Control Protocol
URDB	User Requirements Data Base
WIN	WWMCCS Interconnect Network

DISTRIBUTION LIST

R100/R101 - 1
R103 - 1
R110 - 1
R120 - 1
R141 - 12

R200 - 1
R400 - 1
R600 - 1
R800 - 1
NCS-TS - 1
B101 - 1
308C - 1
308 - 1
J700 - 1

DCA-EUR - 2 (Defense Communications Agency European Area
ATTN: Technical Library
APO New York 09131)

DCA-PAC - 3 (Commander
Defense Communications Agency Pacific Area
Wheeler AFB, HI 96854)

DCA SW PAC - 1 (Commander, DCA - Southwest Pacific Region
APO San Francisco 96274)

DCA NW PAC - 1 (Commander, DCA - Northwest Pacific Region
APO San Francisco 96301)

DCA KOREA - 1 (Chief, DCA - Korea Field Office
APO San Francisco 96301)

DCA-Okinawa - 1 (Chief, DCA - Okinawa Field Office
FPO Seattle 98773)

DCA-Guam - 1 (Chief, DCA - Guam Field Office
Box 141 NAVCAMS WESTPAC
FPO San Francisco 96630)

US NAV Shore EE PAC - 1 (U.S. Naval Shore Electronics Engineering
Activity Pacific, Box 130, ATTN: Code 420
Pearl Harbor, HI 96860)

USDCFO - 1 (Unclassified/Unlimited Distribution)
(Chief, USDCFO/US NATO
APO New York 09667)

END

FILMED

11-83

DETIC